MPLS Working Group   B. Niven-Jenkins, Ed.
Internet-Draft   BT
Intended status: Standards Track   D. Brungard, Ed.
Expires: October 6, 2009   AT&T
                                                            M. Betts, Ed.
                                                         Nortel Networks
                                                             N. Sprecher
                                                  Nokia Siemens Networks
                                                                 S. Ueno
                                                                     NTT
                                                           April 4, 2009

MPLS-TP Requirements
draft-ietf-mpls-tp-requirements-06

Status of this Memo

Copyright Notice

Abstract

This document specifies the requirements of an MPLS Transport Profile (MPLS-TP). This document is a product of a joint International Telecommunications Union (ITU)-IETF effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by International Telecommunications Union - Telecommunications Standardization Sector (ITU-T).

This work is based on two sources of requirements; MPLS and PWE3 architectures as defined by IETF, and packet transport networks as defined by ITU-T.

The requirements expressed in this document are for the behavior of the protocol mechanisms and procedures that constitute building blocks out of which the MPLS transport profile is constructed. The requirements are not implementation requirements.

Requirements Language

Although this document is not a protocol specification, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] and are to be interpreted as instructions to the protocol designers producing solutions that satisfy the requirements set out in this document.

Table of Contents

1. Introduction 4
  1.1. Terminology 6
    1.1.1. Abbreviations 6
    1.1.2. Definitions 8
  1.2. Transport network overview 11
  1.3. Layer network overview 12
2. MPLS-TP Requirements 13

2.1. General requirements 13
2.2. Layering requirements 16
2.3. Data plane requirements 17
2.4. Control plane requirements 18
2.5. Network Management (NM) requirements 20
2.6. Operation, Administration and Maintenance (OAM) requirements 20
2.7. Network performance management (PM) requirements 20
2.8. Recovery requirements 20
2.8.1. Data plane behavior requirements 21
2.8.1.1. Protection 21
2.8.1.2. Restoration 22
2.8.1.3. Sharing of protection resources 22
2.8.1.4. Reversion 23
2.8.2. Triggers for protection, restoration, and reversion . 23
2.8.3. Management plane operation of protection and restoration 23
2.8.4. Control plane and in-band OAM operation of recovery . 24
2.8.5. Topology-specific recovery mechanisms 25
2.8.5.1. Ring protection 25
2.9. QoS requirements 28
2.10. Security requirements 28
3. IANA Considerations 29
4. Security Considerations 29
5. Acknowledgements 29
6. References 29
6.1. Normative References 29
6.2. Informative References 30
Authors' Addresses 31

1. Introduction

Bandwidth demand continues to grow worldwide, stimulated by the accelerating growth and penetration of new packet based services and multimedia applications:

- Packet-based services such as Ethernet, Voice over IP (VoIP), Layer 2 (L2)/Layer 3 (L3) Virtual Private Networks (VPNs), IP Television (IPTV), Radio Access Network (RAN) backhauling, etc.,

- Applications with various bandwidth and Quality of Service (QoS) requirements.

This growth in demand has resulted in dramatic increases in access rates that are, in turn, driving dramatic increases in metro and core network bandwidth requirements.

Over the past two decades, the evolving optical transport infrastructure (Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH), Optical Transport Network (OTN)) has provided carriers with a high benchmark for reliability and operational simplicity.

With the movement towards packet based services, the transport network has to evolve to encompass the provision of packet aware capabilities while enabling carriers to leverage their installed, as well as planned, transport infrastructure investments.

Carriers are in need of technologies capable of efficiently supporting packet based services and applications on their transport networks with guaranteed Service Level Agreements (SLAs). The need to increase their revenue while remaining competitive forces operators to look for the lowest network Total Cost of Ownership (TCO). Investment in equipment and facilities (Capital Expenditure (CAPEX)) and Operational Expenditure (OPEX) should be minimized.

There are a number of technology options for carriers to meet the challenge of increased service sophistication and transport efficiency, with increasing usage of hybrid packet transport and circuit transport technology solutions. To realize these goals, it is essential that packet transport technology be available that can support the same high benchmarks for reliability and operational simplicity set by SDH/SONET and OTN technologies.

Furthermore for carriers it is important that operation of such packet transport networks should preserve the look-and-feel to which carriers have become accustomed in deploying their optical transport networks, while providing common, multi-layer operations, resiliency,

control and multi-technology management.

Transport carriers require control and deterministic usage of network resources. They need end-to-end control to engineer network paths and to efficiently utilize network resources. They require capabilities to support static (management plane based) or dynamic (control plane based) provisioning of deterministic, protected and secured services and their associated resources.

It is also important to ensure smooth interworking of the packet transport network with other existing/legacy packet networks, and provide mappings to enable packet transport carriage over a variety of transport network infrastructures. The latter has been termed vertical interworking, and is also known as client/server or network interworking. The former has been termed horizontal interworking, and is also known as peer-partition or service interworking. For more details on interworking and some of the issues that may arise

<<Comment #1 (Editorial)>>
OLD

(especially with horizontal interworking) seeG.805 [ITU.G805.2000]
NEW
(especially with horizontal interworking) see G.805 [ITU.G805.2000]
<<End Comment>>
and Y.1401 [ITU.Y1401.2008].

Multi-Protocol Label Switching (MPLS) is a maturing packet technology and it is already playing an important role in transport networks and services. However, not all of MPLS's capabilities and mechanisms are needed and/or consistent with transport network operations. There are also transport technology characteristics that are not currently reflected in MPLS. There is therefore the need to define an MPLS Transport Profile (MPLS-TP) that supports the capabilities and functionalities needed for packet transport network services and operations through combining the packet experience of MPLS with the operational experience and practices of existing transport networks.

<<Comment #2>>
The paragraph below was proposed to be removed via a WG LC comment. If it is not possible to remove, the following rephrase is proposed.
OLD
MPLS-TP will enable the migration of transport networks to a packet-based network that will efficiently scale to support packet services in a simple and cost effective way. MPLS-TP needs to combine the necessary existing capabilities of MPLS with additional minimal mechanisms in order that it can be used in a transport role.
NEW
MPLS-TP will enable the ~~migration of transport networks to a~~deployment of packet-based networks that will efficiently scale to support packet services in a simple and cost effective way. MPLS-TP needs to combine the necessary existing capabilities of MPLS with additional minimal mechanisms in order that it can be used in a transport role.
<<End Comment>>

This document specifies the requirements of an MPLS Transport Profile (MPLS-TP). The requirements are for the behavior of the protocol mechanisms and procedures that constitute building blocks out of which the MPLS transport profile is constructed. That is, the requirements indicate what features are to be available in the MPLS toolkit for use by MPLS-TP. The requirements in this document do not describe what functions an MPLS-TP implementation supports. The purpose of this document is to identify the toolkit and any new protocol work that is required.

Although this document is not a protocol specification, the key words

"MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are used as described in [RFC2119] and are to be interpreted as instructions to the protocol designers producing solutions that satisfy the requirements set out in this document.

This document is a product of a joint ITU-T and IETF effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by ITU-T.

This work is based on two sources of requirements, MPLS and PWE3 architectures as defined by IETF and packet transport networks as defined by ITU-T. The requirements of MPLS-TP are provided below. The relevant functions of MPLS and PWE3 are included in MPLS-TP, except where explicitly excluded.

Although both static and dynamic configuration of MPLS-TP transport paths (including Operations, Administration and Maintenance (OAM) and protection capabilities) is required by this document, it MUST be possible for operators to be able to completely operate (including OAM and protection capabilities) an MPLS-TP network in the absence of

<<Comment #3>>
The following rephrase was proposed via a WG LC comment.
As it was not accepted, it is not clear what is the different between “control plane protocols” and “control plane protocols for dynamic configuration”.

OLD

any control plane protocols for dynamic configuration.

NEW

any control plane protocols~~ for dynamic configuration~~.

<<End Comment>>

1.1. Terminology

<<Comment #4 (Editorial)>>
It is proposed to explicitly reference the Rosetta stone draft (informatively):

OLD

Note: Mapping between the terms in this section and ITU-T terminology will be described in a subsequent document.

NEW

Note: Mapping between the terms in this section and ITU-T terminology ~~will be~~is described in [I-D. helvoort-mpls-tp-rosetta-stone]~~a subsequent document~~.

<<End Comment>>

The recovery requirements in this document use the recovery terminology defined in RFC 4427 [RFC4427], this is applied to both control plane and management plane based operations of MPLS-TP transport paths.

1.1.1. Abbreviations

ASON: Automatically Switched Optical Network

ASTN: Automatic Switched Transport Network

ATM: Asynchronous Transfer Mode

CAPEX: Capital Expenditure

CE: Customer Edge

FR: Frame Relay

GMPLS: Generalised Multi-Protocol Label Switching

IGP: Interior Gateway Protocol

IPTV: IP Television

L2: Layer 2

L3: Layer 3

LSP: Label Switched Path

LSR: Label Switching Router

MEP: Maintenance End Point

MIP: Maintenance Intermediate Point

MPLS: Multi-Protocol Label Switching

OAM: Operations, Administration and Maintenance

OPEX: Operational Expenditure

OSI: Open Systems Interconnection

OTN: Optical Transport Network

P2MP: Point to Multi-Point

P2P: Point to Point

PDU: Protocol Data Unit

PM: Performance Management

PSC: Protection State Coordination

PW: Pseudo Wire

QoS: Quality of Service

RAN: Radio Access Network

SDH: Synchronous Digital Hierarchy

SLA: Service Level Agreement

SLS: Service Level Specification

S-PE: Switching Provider Edge

SONET: Synchronous Optical Network

SRLG: Shared Risk Link Group

TCO: Total Cost of Ownership

T-PE: Terminating Provider Edge

VoIP: Voice over IP

VPN: Virtual Private Network

WDM: Wavelength Division Multiplexing

1.1.2. Definitions

Note: The definition of segment in a GMPLS/ASON context (i.e. as defined in RFC4397 [RFC4397]) encompasses both segment and concatenated segment as defined in this document.

Associated bidirectional path: A path that supports traffic flow in both directions but which is constructed from a pair of unidirectional paths (one for each direction) which are associated with one another at the path's ingress/egress points. The forward and backward directions are setup, monitored and protected independently. As a consequence they may or may not follow the same route (links and nodes) across the network.

Client layer network: In a client/server relationship (see G.805 [ITU.G805.2000]), the client layer network receives a (transport) service from the lower server layer network (usually the layer network under consideration).

Concatenated Segment: A serial-compound link connection as defined in G.805 [ITU.G805.2000]. A concatenated segment is a contiguous part of an LSP or multi-segment PW that comprises a set of segments and their interconnecting nodes in sequence. See also "Segment".

Co-routed Bidirectional path: A path where the forward and backward directions follow the same route (links and nodes) across the network. Both directions are setup, monitored and protected as a single entity.

Domain: A domain represents a collection of entities (for example

network elements) that are grouped for a particular purpose, examples of which are administrative and/or managerial responsibilities, trust relationships, addressing schemes, infrastructure capabilities, aggregation, survivability techniques, distributions of control functionality, etc. Examples of such domains include IGP areas and Autonomous Systems.

Layer network: Layer network is defined in G.805 [ITU.G805.2000]. A layer network provides for the transfer of client information and independent operation of the client OAM. A Layer Network may be described in a service context as follows: one layer network may provide a (transport) service to higher client layer network and may, in turn, be a client to a lower layer network. A layer network is a logical construction somewhat independent of arrangement or composition of physical network elements. A particular physical network element may topologically belong to more than one layer network, depending on the actions it takes on the encapsulation associated with the logical layers (e.g. the label stack), and thus could be modeled as multiple logical elements. A layer network may consist of one or more sublayers. Section 1.3 provides a more detailed overview of what constitutes a layer network. For additional explanation of how layer networks relate to the OSI concept of layering see Appendix I of Y.2611 [ITU.Y2611.2006].

Link: A physical or logical connection between a pair of LSRs that are adjacent at the (sub)layer network under consideration. A link may carry zero, one or more LSPs or PWs. A packet entering a link will emerge with the same label stack entry values.

MPLS-TP Logical Ring: An MPLS-TP logical ring is constructed from a set of LSRs and logical data links (such as MPLS-TP LSP tunnels or MPLS-TP pseudowires) and physical data links that form a ring topology.

Path: See Transport Path.

MPLS-TP Physical Ring: An MPLS-TP physical ring is constructed from a set of LSRs and physical data links that form a ring topology.

MPLS-TP Ring Topology: In an MPLS-TP ring topology each LSR is connected to exactly two other LSRs, each via a single point-to-point bidirectional MPLS-TP capable link. A ring may also be constructed from only two LSRs where there are also exactly two links. Rings may be connected to other LSRs to form a larger network. Traffic originating or terminating outside the ring may be carried over the ring. Client network nodes (such as CEs) may be connected directly to an LSR in the ring.

<<Comment #5>>
We think that the definition of Section Layer Network below is not fully in line with the ITU-T definition as given in the LS from last December 2008 meeting. We are planning to provide some text proposal for the Section Layer Network definition to resolve this comment.
<<End Comment>>

Section Layer Network: A section is a server layer (which may be MPLS-TP or a different technology) which provides for encapsulation and OAM of a client layer network. A section layer may provide for aggregation of multiple MPLS-TP clients. Note that G.805 [ITU.G805.2000] defines the section layer as one of the two layer networks in a transmission media layer network. The other layer network is the physical media layer network.

Segment: A link connection as defined in G.805 [ITU.G805.2000]. A segment is the part of an LSP that traverses a single link or the part of a PW that traverses a single link (i.e. that connects a pair of adjacent {Switching|Terminating} Provider Edges). See also "Concatenated Segment".

Server Layer Network: In a client/server relationship (see G.805 [ITU.G805.2000]), the server layer network provides a (transport) service to the higher client layer network (usually the layer network under consideration).

Sublayer: Sublayer is defined in G.805 [ITU.G805.2000]. The distinction between a layer network and a sublayer is that a sublayer is not directly accessible to clients outside of its encapsulating layer network and offers no direct transport service for a higher layer (client) network.

Switching Provider Edge (S-PE): See [I-D.ietf-pwe3-ms-pw-arch].

Tandem Connection: A tandem connection is an arbitrary part of a transport path that can be monitored (via OAM) independently from the end-to-end monitoring (OAM). It may be a monitored segment or a monitored concatenated segment of a transport path. The tandem connection may also include the forwarding engine(s) of the node(s) at the edge(s) of the segment or concatenated segment.

Terminating Provider Edge (T-PE): See [I-D.ietf-pwe3-ms-pw-arch].

Transport Path: A network connection as defined in G.805 [ITU.G805.2000]. In an MPLS-TP environment a transport path corresponds to an LSP or a PW.

<<Comment #6 (Technical)>>
There is a need to discuss the text changes below, proposed via WG LC comments.
OLD

Transport Path Layer: A layer network that provides point-to-point or point-to-multipoint transport paths that may be used to carry a higher (client) layer network or aggregates of higher (client) layer networks, for example the transport service layer. It provides independent (of the client) OAM when transporting its clients.

NEW

Transport Path Layer: A (sub-)layer network that provides point-to-point or point-to-multipoint transport paths that may be used to carry ~~a higher (client) layer network or~~ aggregates of transport paths in the higher (client) layer networks, ~~for example~~that can be the transport service layer or the transport path layer itself~~. It provides independent (of the client) OAM when transporting its clients~~.

<<End Comment>>

Transport Service Layer: A layer network in which transport paths are used to carry a customer's (individual or bundled) service (may be point-to-point, point-to-multipoint or multipoint-to-multipoint services).

Transmission Media Layer: A layer network, consisting of a section layer network and a physical layer network as defined in G.805 [ITU.G805.2000], that provides sections (two-port point-to-point connections) to carry the aggregate of network transport path or network service layers on various physical media.

Unidirectional Path: A path that supports traffic flow in only one direction.

1.2. Transport network overview

The connectivity service is the basic service provided by a transport network. The purpose of a transport network is to carry its customer traffic (i.e. the stream of customer PDUs or customer bits, including overhead) between endpoints in the transport network (typically over several intermediate nodes). The connectivity services offered to customers are typically aggregated into large transport paths with long-holding times and independent OAM (of the client OAM), which contribute to enabling the efficient and reliable operation of the transport network. These transport paths are modified infrequently.

Quality-of-service mechanisms are required in the packet transport network to ensure the prioritization of critical services, to guarantee bandwidth and to control jitter and delay. A transport network must provide the means to commit quality of service objectives to clients. This is achieved by providing a mechanism for client network service demarcation for the network path together with an associated network resiliency mechanism.

Aggregation is beneficial for achieving scalability and security since:

1. It reduces the number of provisioning and forwarding states in the network core.

2. It reduces load and the cost of implementing service assurance and fault management.

3. Customer traffic is encapsulated and layer associated OAM overhead is added. This allows complete isolation of customer traffic and its management from carrier operations.

An important attribute of a transport network is that it is able to function regardless of which clients are using its connection service or over which transmission media it is running. The client,

transport network and server layers are from a functional and operations point of view independent layer networks. Another key characteristic of transport networks is the capability to maintain the integrity of the client across the transport network. A transport network must also provide a method of service monitoring in order to verify the delivery of an agreed quality of service. This is enabled by means of carrier-grade OAM tools.

Customer traffic is first encapsulated within the transport service layer network. The transport service layer network signals may then be aggregated into a transport path layer network for transport through the network in order to optimize network management. Transport service layer network OAM is used to monitor the transport integrity of the customer traffic and transport path layer network OAM is used to monitor the transport integrity of the aggregates. At any hop, the aggregated signals may be further aggregated in lower layer transport network paths for transport across intermediate shared links. The transport service layer network signals are extracted at the edges of aggregation domains, and are either delivered to the customer or forwarded to another domain. In the core of the network, only the transport path layer network signals are monitored at intermediate points; individual transport service layer network signals are monitored at the network boundary. Although the connectivity of the transport service layer network may be point-to-point, point-to-multipoint or multipoint-to-multipoint, the transport path layer network only provides point-to-point or point-to-multipoint transport paths which are used to carry aggregates of transport service layer network traffic.

1.3. Layer network overview

A layer network provides its clients with a transport service and the operation of the layer network is independent of whatever client happens to use the layer network. Information that passes between any client to the layer network is common to all clients and is the minimum needed to be consistent with the definition of the transport service offered. The client layer network can be connectionless, connection oriented packet switched, or circuit switched. The transport service transfers a payload (individual packet payload for connectionless networks, a sequence of packets payloads in the case of connection oriented packet switched networks, and a deterministic schedule of payloads in the case of circuit switched networks) such that the client can populate the payload without affecting any operation within the serving layer network.

The operations within a layer network that are independent of its clients include the control of forwarding, the control of resource reservation, the control of traffic demerging, and the OAM and

recovery of the transport service. All of these operations are internal to a layer network. By definition, a layer network does not rely on any client information to perform these operations and therefore all information required to perform these operations is independent of whatever client is using the layer network.

A layer network will have consistent features in order to support the control of forwarding, resource reservation, OAM and recovery. For example, a layer network will have a common addressing scheme for the end points of the transport service and a common set of transport descriptors for the transport service. However, a client may use a different addressing scheme or different traffic descriptors (consistent with performance inheritance).

It is sometimes useful to independently monitor a smaller domain within a layer network (or the transport services that traverse this smaller domain) but the control of forwarding or the control of resource reservation involved retain their common elements. These smaller monitored domains are sublayers.

It is sometimes useful to independently control forwarding in a smaller domain within a layer network but the control of resource reservation and OAM retain their common elements. These smaller domains are partitions of the layer network.

2. MPLS-TP Requirements

This document specifies the requirements of an MPLS Transport Profile (MPLS-TP). The requirements are for the behavior of the protocol mechanisms and procedures that constitute building blocks out of which the MPLS transport profile is constructed. That is, the requirements indicate what features are to be available in the MPLS toolkit for use by MPLS-TP. The requirements in this document do not describe what functions an MPLS-TP implementation supports. The purpose of this document is to identify the toolkit and any new protocol work that is required.

2.1. General requirements

1 The MPLS-TP data plane MUST be a subset of the MPLS data plane as defined by the IETF. When MPLS offers multiple options in this respect, MPLS-TP SHOULD select the minimum sub-set (necessary and sufficient subset) applicable to a transport network application.

2 Any new functionality that is defined to fulfill the requirements for MPLS-TP MUST be agreed within the IETF through the IETF

<<Comment #7>>
The use of the MUST seems to be in contradiction with “(as far as practically possible)”. The fact that it can be practically not possible to re-use existing standards, qualify the requirement as a SHOULD.
OLD
consensus process and MUST re-use (as far as practically
NEW
consensus process and ~~MUST~~ SHOULD re-use (as far as practically
<<End Comment>>
possible) existing MPLS standards.

3 Mechanisms and capabilities MUST be able to interoperate with existing IETF MPLS [RFC3031] and IETF PWE3 [RFC3985] control and data planes where appropriate.

A. Data plane interoperability MUST NOT require a gateway function.

4 MPLS-TP and its interfaces, both internal and external, MUST be sufficiently well-defined that interworking equipment supplied by multiple vendors will be possible both within a single domain, and between domains.

5 MPLS-TP MUST be a connection-oriented packet switching technology with traffic engineering capabilities that allow deterministic control of the use of network resources.

6 MPLS-TP MUST support traffic engineered point to point (P2P) and point to multipoint (P2MP) transport paths.

7 MPLS-TP MUST support the logical separation of the control and management planes from the data plane.

8 MPLS-TP MUST support the physical separation of the control and management planes from the data plane.

9 MPLS-TP MUST support static provisioning of transport paths via the management plane.

<<Comment #8>>
The objective of similar operations is applicable to different transport layer networks rather than to different networks.
OLD

10 Mechanisms in an MPLS-TP network that satisfy functional requirements that are common to general transport networks (i.e., independent of technology) SHOULD be operable in a way that is similar to the way the equivalent mechanisms are operated in other transport networks.

NEW

10 Mechanisms in an MPLS-TP layer network that satisfy functional requirements that are common to general transport layer networks (i.e., independent of technology) SHOULD be operable in a way that is similar to the way the equivalent mechanisms are operated in other transport layer networks.

<<End Comment>>

11 Static provisioning MUST NOT depend on the presence of any element of a control plane.

12 MPLS-TP MUST support the capability for network operation (including OAM and recovery) via the management plane (without the use of any control plane protocols).

13 A solution MUST be defined to support dynamic provisioning and restoration of MPLS-TP transport paths via a control plane.

14 MPLS-TP MUST support the co-existence of statically and dynamically provisioned/managed MPLS-TP transport paths within the same layer network or domain.

15 The MPLS-TP data plane MUST be capable of

A. forwarding data independent of the control or management plane used to configure and operate the MPLS-TP layer network.

B. taking recovery actions independent of the control plane used to configure the MPLS-TP layer network. If the control plane does not restart, the data plane connections MUST be held and NOT time out.

<<Comment #9>>
The second part of Req.15B is not clear. Does it mean that control plane failures must not impact data plane?
Once the intention is clear, we can provide some text proposal for Req.15 to resolve this comment.
In the past the requirement was saying: "the MPLS-TP data plane MUST continue to operate normally if the management plane or control plane that configured the transport paths fails".
That text was almost ok for us: there was only the need to clarify that data plane means forwarding, OAM and protection.
<<End Comment>>

16 MPLS-TP MUST support mechanisms to avoid or minimize traffic impact (e.g. packet delay, reordering and loss) during network reconfiguration.

17 MPLS-TP MUST support transport paths through multiple homogeneous domains.

18 MPLS-TP SHOULD support transport paths through multiple non-homogeneous domains.

19 MPLS-TP MUST NOT dictate the deployment of any particular network topology either physical or logical, however:

A. It MUST be possible to deploy MPLS-TP in rings.

B. It MUST be possible to deploy MPLS-TP in arbitrarily interconnected rings with one or two points of interconnection.

C. MPLS-TP MUST support rings of at least 16 nodes in order to support the upgrade of existing TDM rings to MPLS-TP. MPLS-TP SHOULD support rings with more than 16 nodes.

20 MPLS-TP MUST be able to scale at least as well as existing transport technologies with growing and increasingly complex network topologies as well as with increasing bandwidth demands, number of customers, and number of services.

21 MPLS-TP SHOULD support mechanisms to safeguard against the provisioning of transport paths which contain forwarding loops.

2.2. Layering requirements

22 A generic and extensible solution MUST be provided to support the transport of one or more client layer networks (e.g. MPLS-TP, IP, MPLS, Ethernet, ATM, FR, etc.) over an MPLS-TP layer network.

23 A generic and extensible solution MUST be provided to support the transport of MPLS-TP transport paths over one or more server layer networks (such as MPLS-TP, Ethernet, SONET/SDH, OTN, etc.). Requirements for bandwidth management within a server layer network are outside the scope of this document.

24 In an environment where an MPLS-TP layer network is supporting a client layer network, and the MPLS-TP layer network is supported by a server layer network then operation of the MPLS-TP layer network MUST be possible without any dependencies on the server or client layer network.

A. The server layer MUST guarantee that the traffic loading imposed by other clients does not cause the transport service provided to the MPLS-TP layer to fall bellow the agreed level. Mechanisms to achieve this are outside the scope of these requirements.

B. It MUST be possible to isolate the control and management planes of the MPLS-TP layer network from the control and management planes of the client and server layer networks.

<<Comment #10>>
We think that the option to have unified control/management of multi-layer networks encompassing MPLS-TP should be allowed. Please add Req.24C:
"
C. In case client and server layer networks are located in the same administrative domain, it MUST be possible to operate client and server layer networks via a common control or management plane.
"
<<End Comment>>

25 A solution MUST be provided to support the transport of a client MPLS or MPLS-TP layer network over a server MPLS or MPLS-TP layer network.

A. The level of co-ordination required between the client and server MPLS(-TP) layer networks MUST be minimized (preferably no co-ordination will be required).

B. The MPLS(-TP) server layer network MUST be capable of transporting the complete set of packets generated by the client MPLS(-TP) layer network, which may contain packets that are not MPLS packets (e.g. IP or CNLS packets used by the control/management plane of the client MPLS(-TP) layer network).

26 It MUST be possible to operate the layers of a multi-layer network that includes an MPLS-TP layer autonomously.

The above are not only technology requirements, but also operational requirements. Different administrative groups may be responsible for the same layer network or different layer networks.

27 It MUST be possible to hide MPLS-TP layer network addressing and other information (e.g. topology) from client layer networks. However, it SHOULD be possible, at the option of the operator, to leak a limited amount of summarized information (such as SRLGs or reachability) between layers.

2.3. Data plane requirements

28 It MUST be possible for the end points of an MPLS-TP transport path that is carrying an aggregate of client transport paths to be able to decompose the aggregate transport path into its component client transport paths.

29 A transport path on a link MUST be uniquely identifiable by a single label on that link.

30 A transport path's source MUST be identifiable at its destination within its layer network in coordination with the management plane or control plane.

<<Comment #11>>
It is not clear what is the meaning of “in coordination with the management plane or control plane” in Req.30.
Once the intention is clear, we can provide some text proposal for Req.30 to resolve this comment.
<<End Comment>>

31 MPLS-TP MUST be capable of using P2MP server (sub-)layer capabilities as well as P2P server (sub-)layer capabilities when supporting P2MP MPLS-TP transport paths.

<<Comment #12>>
In transport networks, only unidirectional and co-routed bidirectional paths are required. As a consequence, associated bidirectional paths are not required in MPLS-TP.
If there is a requirement to extend the MPLS toolkit to support associated bidirectional paths as part of the MPLS-TP standardization efforts, it is proposed to clarify in Req.32 below that associated bidirectional paths are required for MPLS and not MPLS-TP environments.
We are planning to provide some text proposal for Req.31 to resolve this comment.
<<End Comment>>

32 MPLS-TP MUST support unidirectional, co-routed bidirectional and associated bidirectional point-to-point transport paths.

<<Comment #xx>>
Propose rephrase for Req.33, 34, 35 and 36
<<End Comment>>

33 The end points of a co-routed bidirectional transport path MUST be aware of the pairing relationship of the forward and reverse paths used to support the bidirectional service.

34 The intermediate nodes (including MEPs, MIPs and other internal functions as appropriate) of a co-routed bidirectional transport path at each (sub-)layer MUST be aware of the pairing relationship of the forward and the backward directions of that transport path.

35 The end points of an associated bidirectional transport path MUST be aware of the pairing relationship of the forward and reverse paths used to support the bidirectional service.

36 The intermediate nodes (including MEPs, MIPs and other internal functions as appropriate) of an associated bidirectional transport path at each (sub-)layer SHOULD NOT be aware of the pairing relationship of the forward and the backward directions of that transport path.

37 MPLS-TP MUST support bidirectional transport paths with asymmetric bandwidth requirements, i.e. the amount of reserved bandwidth differs between the forward and backward directions.

38 MPLS-TP MUST support unidirectional point-to-multipoint transport paths.

39 MPLS-TP MUST be extensible in order to accommodate new types of client layer networks and services.

40 MPLS-TP SHOULD support mechanisms to enable the reserved bandwidth associated with a transport path to be increased without impacting the existing traffic on that transport path provided enough resources are available.

41 MPLS-TP SHOULD support mechanisms to enable the reserved bandwidth of a transport path to be decreased without impacting the existing traffic on that transport path, provided that the level of existing traffic is smaller than the reserved bandwidth following the decrease.

42 MPLS-TP MUST support mechanisms which ensure the integrity of the transported customer's service traffic as required by its associated SLA. Loss of integrity may be defined as packet corruption, re-ordering or loss during normal network conditions.

43 MPLS-TP MUST support mechanisms to detect when loss of integrity of the transported customer's service traffic has occurred.

44 MPLS-TP MUST support an unambiguous and reliable means of distinguishing users' (client) packets from MPLS-TP control packets (e.g. control plane, management plane, OAM and protection switching packets).

2.4. Control plane requirements

This section defines the requirements that apply to an MPLS-TP control plane. Note that it MUST be possible to operate an MPLS-TP network without using a control plane.

<<Comment #13>>
The WG LC comment regarding ASTN and the references for ASON architecture and requirements needs further discussion.
<<End Comment>>

The ITU-T has defined an architecture for Automatically Switched Optical and Transport Networks (ASON/ASTN) in G.8080 [ITU.G8080.2006]

and G.8080 Amd1 [ITU.G8080.2008]. The control plane for MPLS-TP MUST fit within the ASON/ASTN architecture.

An interpretation of the ASON/ASTN signaling and routing requirements in the context of GMPLS can be found in [RFC4139] and [RFC4258].

Additionally:

<<Comment #14>>
There is a need to discuss the following WG LC comment. Moreover the requirement looks to be generic rather than just for the control plane.
OLD

45 It MUST be possible to operate and configure the MPLS-TP data plane without any IP forwarding capability in the MPLS-TP data plane.

NEW

45 It MUST be possible to operate and configure the MPLS-TP data plane (forwarding, OAM and protection) without any IP ~~forwarding~~ capability~~ in the MPLS-TP data plane~~.

<<End Comment>>

46 The MPLS-TP control pane MUST support control plane topology and data plane topology independence. As a consequence a failure of the control plane does not imply that there has also been a failure of the data plane.

47 The MPLS-TP control plane MUST be able to be operated independent of any particular client or server layer control plane.

48 MPLS-TP SHOULD define a solution to support an integrated control plane encompassing MPLS-TP together with its server and client

layer networks when these layer networks belong to the same administrative domain.

49 The MPLS-TP control plane MUST support establishing all the connectivity patterns defined for the MPLS-TP data plane (e.g.,

<<Comment #15>>
For clarify, it might be worth spelling both associated and co-routed bidirectional paths:
OLD
unidirectional and bidirectional P2P, unidirectional P2MP, etc.)
NEW
Unidirectional and associated and co-routed bidirectional P2P, unidirectional P2MP, etc.)
<<End Comment>>

including configuration of protection functions and any associated maintenance functions.

50 The MPLS-TP control plane MUST support the configuration and modification of OAM maintenance points as well as the activation/ deactivation of OAM when the transport path or transport service is established or modified.

51 An MPLS-TP control plane MUST support operation of the recovery functions described in Section 2.8.

52 An MPLS-TP control plane MUST scale gracefully to support a large number of transport paths, nodes and links.

<<Comment #16>>
It is proposed to rephrase Req.53 as a requirement for the solution and not for an implementation (given the scope of this document).
OLD
53 If a control plane is used for MPLS-TP, the control plane's graceful restart capabilities, if any, MUST be supported.
NEW
53 If a control plane is used for MPLS-TP, the control plane~~'s~~ MUST be capable of gracefully restart~~ capabilities, if any, MUST be supported~~.
<<End Comment>>

54 An MPLS-TP control plane MUST provide a mechanism for dynamic ownership transfer of the control of MPLS-TP transport paths from the management plane to the control plane and vice versa. The number of reconfigurations required in the data plane MUST be minimized (preferably no data plane reconfiguration will be required).

2.5. Network Management (NM) requirements

For requirements related to NM functionality (Management Plane in ITU-T terminology) for MPLS-TP, see the MPLS-TP NM requirements document [I-D.ietf-mpls-tp-nm-req].

2.6. Operation, Administration and Maintenance (OAM) requirements

For requirements related to OAM functionality for MPLS-TP, see the MPLS-TP OAM requirements document [I-D.ietf-mpls-tp-oam-requirements].

2.7. Network performance management (PM) requirements

For requirements related to PM functionality for MPLS-TP, see the MPLS-TP OAM requirements document [I-D.ietf-mpls-tp-oam-requirements].

2.8. Recovery requirements

Network survivability plays a critical role in the delivery of reliable services. Network availability is a significant contributor to revenue and profit. Service guarantees in the form of SLAs require a resilient network that rapidly detects facility or node failures and restores network operation in accordance with the terms of the SLA.

55 MPLS-TP MUST provide protection and restoration mechanisms.

   A. MPLS-TP recovery techniques SHOULD be identical (or as similar as possible) to those already used in existing transport networks to simplify implementation and operations. However, this MUST NOT override any other requirement.

   B. Recovery techniques used for P2P and P2MP SHOULD be identical to simplify implementation and operation. However, this MUST NOT override any other requirement.

56 MPLS-TP recovery mechanisms MUST be applicable at various levels throughout the network including support for link, transport path, segment, concatenated segment and end to end recovery.

57 MPLS-TP recovery paths MUST meet the SLA protection objectives of the service.

   A. MPLS-TP MUST provide mechanisms to guarantee 50ms recovery times from the moment of fault detection in networks with spans less than 1200 km.

   B. For protection it MUST be possible to require protection of 100% of the traffic on the protected path.

   C. Recovery objectives SHOULD be configurable per transport path, and SHOULD support objectives for bandwidth and QoS.

<<Comment #17>>
Req.57C is still not fully clear.
Once the intention is clear, we can provide some text proposal for Req.57C to resolve this comment.
<<End Comment>>

D. Recovery MUST meet SLA requirements over multiple domains.

58 The recovery mechanisms SHOULD be applicable to any topology.

59 The recovery mechanisms MUST support the means to operate in synergy with (including coordination of timing) the recovery mechanisms present in any client or server transport networks (for example, Ethernet, SDH, OTN, WDM) to avoid race conditions between the layers.

60 MPLS-TP recovery and reversion mechanisms MUST prevent frequent operation of recovery in the event of an intermittent defect.

2.8.1. Data plane behavior requirements

General protection and survivability requirements are expressed in terms of the behavior in the data plane.

2.8.1.1. Protection

Note: Only nodes that are aware of the pairing relationship between the forward and backward directions of an associated bidirectional transport path can be used as end points to protect all or part of that transport path.

61 MPLS-TP MUST support 1+1 protection.

A. Bidirectional 1+1 protection for P2P connectivity MUST be supported.

B. Unidirectional 1+1 protection for P2P connectivity MUST be supported.

C. Unidirectional 1+1 protection for P2MP connectivity MUST be supported.

62 MPLS-TP MUST support 1:n protection (including 1:1 protection).

<<Comment #18>>
The WG LC comment regarding the difference between 1:n and (1:1)^n needs further discussion
<<End Comment>>

<<Comment #19 (Editorial)>>

I like the way Req.61A has been rephrased. I would rephrase also Req.62A in the same way:
OLD

A. MPLS-TP 1:n protection MUST include bidirectional protection switching for P2P connectivity, and this SHOULD be the default behavior for 1:n protection.

NEW

A. ~~MPLS-TP~~Bidirectional 1:n protection MUST ~~include bidirectional protection~~ ~~switching~~ for P2P connectivity MUST be supported, and ~~this~~ SHOULD be the default behavior for 1:n protection.

<<End Comment>>

B. Unidirectional 1:n protection for P2MP connectivity MUST be supported.

C. Unidirectional 1:n protection for P2P connectivity is not required and MAY be omitted from the MPLS-TP specifications.

D. The action of protection switching MUST NOT cause user data to loop. Backtracking is allowed.

<<Comment #20>>
It is proposed to add the following requirement that was present in the LS ITU-T has sent to IETF in September 2006:
"
MPLS-TP protection mechanisms MUST be capable to operate without any IP capability.
"
<<End Comment>>

Note: Support for extra traffic (as defined in [RFC4427]) is not required in MPLS-TP and MAY be omitted from the MPLS-TP specifications.

2.8.1.2. Restoration

<<Comment #21>>
The requirements below are not requirements for a data plane behavior. We do not see any requirement on the data plane behavior (other than general MPLS-TP data plane requirements) to support restoration mechanisms. The requirements below apply to the either the control plane or the management plane.
<<End Comment>>

63 The restoration transport path MUST be able to share resources with the transport path being replaced (sometimes known as soft rerouting).

64 Restoration priority MUST be supported so that an implementation can determine the order in which transport paths should be restored (to minimize service restoration time as well as to gain access to available spare capacity on the best paths).

65 Preemption priority MUST be supported to allow restoration to displace other transport paths in the event of resource constraint.

2.8.1.3. Sharing of protection resources

<<Comment #22>>
OLD

66 MPLS-TP SHOULD support 1:n (including 1:1) shared mesh restoration.

NEW

66 MPLS-TP SHOULD support 1:n (including 1:1) shared mesh ~~restoration~~recovery.

<<End Comment>>

67 MPLS-TP MUST support the definition of shared protection groups to allow the coordination of protection actions resulting from triggers caused by events at different locations in the network.

<<Comment #23>>

The definition of “shared protection groups” is missing. As a consequence, Req.67 is not clear.

Once the intention is clear, we can provide some text proposal for Req.67 to resolve this comment.

<<End Comment>>

68 MPLS-TP MUST support sharing of protection resources such that protection paths that are known not to be required concurrently can share the same resources.

2.8.1.4. Reversion

69 MPLS-TP protection mechanisms MUST support revertive and non-revertive behavior.

70 MPLS-TP restoration mechanisms MUST support revertive and non-revertive behavior.

<<Comment #24>>

Please add Req.69A:

“

A. When MPLS-TP restoration is non-revertive, the working transport path MUST be preserved.

“

<<End Comment>>

2.8.2. Triggers for protection, restoration, and reversion

Recovery actions may be triggered from different places as follows:

71 MPLS-TP MUST support physical layer fault indication triggers.

72 MPLS-TP MUST support OAM-based triggers.

73 MPLS-TP MUST support management plane triggers (e.g., forced switch, etc.).

74 There MUST be a mechanism to allow administrative recovery actions to be distinguished from recovery actions initiated by

other triggers.

75 Where a control plane is present, MPLS-TP SHOULD support control plane triggers.

<<Comment #25>>
A clear way to distinguish when control plane protects from when control plane notifies an action performed by data plane, should also be defined.
We are planning to provide some text proposal for Req.75 to resolve this comment.
<<End Comment>>

76 MPLS-TP protection mechanisms MUST support priority logic to negotiate and accommodate coexisting requests (i.e., multiple requests) for protection switching (e.g., administrative requests and requests due to link/node failures).

<<Comment #26>>
We think that there is a need to specify in this document the priority logic that needs to be supported.
We are planning to provide some text proposal for Req.76 to resolve this comment.
<<End Comment>>

2.8.3. Management plane operation of protection and restoration

All functions described here are for control by the operator.

77 It MUST be possible to configure protection paths and protection-to-working path relationships (sometimes known as protection groups).

78 There MUST be support for pre-calculation of recovery paths.

79 There MUST be support for pre-provisioning of recovery paths.

80 The external controls as defined in [RFC4427] MUST be supported.

<<Comment #26>>
The external controls defined in RFC 4427 are not fully aligned with ITU-T requirements for external controls. It is proposed to align Req.73 with ITU-T requirements for external controls.
We are planning to provide some text proposal for Req.80 to resolve this comment.
<<End Comment>>

A. External controls overruled by higher priority requests (e.g., administrative requests and requests due to link/node failures) or unable to be signaled to the remote end (e.g. because of a protection state coordination fail) MUST be dropped.

81 There MUST be support for the configuration of timers used for

recovery operation.

<<Comment #27>>
Not clear the need/role of timers for recovery operations in Req.81. If the objective is to require the configurability of wait-time-to-restore and hold-off timers, it is proposed to spell these out.
Once the intention is clarify, we can provide some text proposal for Req.81 to resolve this comment.
<<End Comment>>

82 Restoration resources MAY be pre-planned and selected a priori, or computed after failure occurrence.

83 When preemption is supported for restoration purposes, it MUST be possible for the operator to configure it.

84 The management plane MUST provide indications of protection events and triggers.

85 The management plane MUST allow the current protection status of all transport paths to be determined.

2.8.4. Control plane and in-band OAM operation of recovery

86 The MPLS-TP control plane (which is not mandatory in an MPLS-TP implementation) MUST be capable of supporting:

A. establishment and maintenance of all recovery entities and functions

B. signaling of administrative control

C. protection state coordination (PSC). Since control plane network topology is independent from the data plane network topology, the PSC supported by the MPLS-TP control plane MAY run on resources different than the data plane resources handled within the recovery mechanism (e.g. backup).

87 In-band OAM MUST be capable of supporting:

A. signaling of administrative control

B. protection state coordination (PSC). Since in-band OAM tools share the data plane with the carried transport service, in order to optimize the usage of network resources, the PSC supported by in-band OAM MUST run on protection resources.

2.8.5. Topology-specific recovery mechanisms

88 MPLS-TP MAY support recovery mechanisms that are optimized for specific network topologies. These mechanisms MUST be interoperable with the mechanisms defined for arbitrary topology

(mesh) networks to enable protection of end-to-end transport paths.

2.8.5.1. Ring protection

Several service providers have expressed a high level of interest in operating MPLS-TP in ring topologies and require a high level of survivability function in these topologies. The requirements listed below have been collected from these service providers and from the ITU-T.

The main objective in considering a specific topology (such as a ring) is to determine whether it is possible to optimize any

<<Comment #28>>
The last sentence is a vendor/operator issue and not a standard/requirement issue. It is proposed to remove it.

OLD

mechanisms such that the performance of those mechanisms within the topology is significantly better than the performance of the generic mechanisms in the same topology. The benefits of such optimizations are traded against the costs of developing, implementing, deploying, and operating the additional optimized mechanisms noting that the generic mechanisms MUST continue to be supported.

NEW

mechanisms such that the performance of those mechanisms within the topology is significantly better than the performance of the generic mechanisms in the same topology~~. The benefits of such optimizations are traded against the costs of developing, implementing, deploying, and operating the additional optimized mechanisms noting that the generic mechanisms MUST continue to be supported~~.

<<End Comment>>

Within the context of recovery in MPLS-TP networks, the optimization criteria considered in ring topologies are as follows:

a. Minimize the number of OAM entities that are needed to trigger the recovery operation - less than are required by other recovery mechanisms.

b. Minimize the number of elements of recovery in the ring - less than are required by other recovery mechanisms.

c. Minimize the number of labels required for the protection paths across the ring - less than are required by other recovery mechanisms.

d. Minimize the amount of control and management plane transactions during a maintenance operation (e.g., ring upgrade) - less than are required by other recovery mechanisms.

e. When a control plane is supported, minimize the impact on signaling and routing information exchange during protection - less than are required by other recovery mechanisms.

<<Comment #29>>
The statement below is not 100% accurate:
- Req.89 is specific to ring topologies
- Req.93 and 101 specifies some behavior that is required for an optimized solution

It is proposed the following rephrase:

OLD

It may be observed that the requirements in this section are fully compatible with the generic requirements expressed above, and that no requirements that are specific to ring topologies have been identified.

NEW

It may be observed that the requirements in this section are fully compatible with the generic requirements expressed above~~, and that no requirements that are specific to ring topologies have been identified~~.

<<End Comment>>

89 MPLS-TP MUST include recovery mechanisms that operate in any single ring supported in MPLS-TP, and continue to operate within the single rings even when the rings are interconnected.

90 When a network is constructed from interconnected rings, MPLS-TP MUST support recovery mechanisms that protect user data that traverses more than one ring. This includes the possibility of failure of the ring-interconnect nodes and links.

91 MPLS-TP recovery in a ring MUST protect unidirectional and bidirectional P2P transport paths.

92 MPLS-TP recovery in a ring MUST protect unidirectional P2MP transport paths.

<<Comment #30>>

OLD

93 MPLS-TP 1+1 and 1:1 protection in a ring MUST support switching time within 50 ms from the moment of fault detection in a network with a 16 nodes ring with less than 1200 km of fiber.

NEW

93 An optimized solution for MPLS-TP protection in a ring MUST support only ~~1+1 and~~ 1:1 protection ~~in a ring MUST support~~ with switching time within 50 ms from the moment of fault detection in a network with a 16 nodes ring with less than 1200 km of fiber.

<<End Comment>>

94 The protection switching time in a ring MUST be independent of the number of LSPs crossing the ring.

95 Recovery actions in a ring MUST be data plane functions triggered by different elements of control. The triggers are configured by management or control planes and are subject to configurable policy.

<<Comment #31>>
Req.95 needs to be clarified.
Once the intention is clarified, we can provide some text proposal for Req.95 to resolve this comment.
<<End Comment>>

96 The configuration and operation of recovery mechanisms in a ring MUST scale well with:

A. the number of transport paths (must be better than linear scaling)

B. the number of nodes on the ring (must be at least as good as linear scaling)

C. the number of ring interconnects (must be at least as good as linear scaling)

97 Recovery techniques used in a ring MUST NOT prevent the ring from being connected to a general MPLS-TP network in any arbitrary way, and MUST NOT prevent the operation of recovery techniques in the rest of the network.

98 MPLS-TP Recovery mechanisms applicable to a ring MUST be equally applicable in physical and logical rings.

99 Recovery techniques in a ring SHOULD be identical (or as similar as possible) to those in general transport networks to simplify implementation and operations. However, this MUST NOT override any other requirement.

100 Recovery techniques in logical and physical rings SHOULD be identical to simplify implementation and operation. However, this MUST NOT override any other requirement.

<<Comment #32>>
OLD
101 The default recovery scheme in a ring MUST be bidirectional recovery in order to simplify the recovery operation.
NEW
101 ~~The default recovery scheme in a ring~~An optimized solution for MPLS-TP protection in a ring MUST ~~be~~ support only bidirectional recovery in order to simplify the recovery operation.
<<End Comment>>

102 The recovery mechanism in a ring MUST support revertive switching, which MUST be the default behavior. This allows optimization of the use of the ring resources, and restores the preferred quality conditions for normal traffic (e.g., delay) when the recovery mechanism is no longer needed.

103 The recovery mechanisms in a ring MUST support ways to allow administrative protection switching, to be distinguished from protection switching initiated by other triggers.

104 It MUST be possible to lockout (disable) protection mechanisms on selected links (spans) in a ring (depending on operator's need). This may require lockout mechanisms to be applied to intermediate nodes within a transport path.

105 MPLS-TP recovery mechanisms in a ring:

A. MUST include a mechanism to allow an implementation to handle (including the coordination of) coexisting requests or triggers (i.e., multiple requests - not necessarily arriving simultaneously and located anywhere in the ring) for protection switching based on priority. Note that such coordination is the ring equivalent of the definition of shared protection groups.

B. MAY support multiple failures without reconfiguring the protection actions.

106 MPLS-TP recovery and reversion mechanisms in a ring MUST offer a way to prevent frequent operation of recovery in the event of an intermittent defect.

107 MPLS-TP MUST support the sharing of protection bandwidth in a ring by allowing best effort traffic.

108 MPLS-TP MUST support sharing of ring protection resources such that protection paths that are known not to be required concurrently can share the same resources.

2.9. QoS requirements

Carriers require advanced traffic management capabilities to enforce and guarantee the QoS parameters of customers' SLAs.

Quality of service mechanisms are REQUIRED in an MPLS-TP network to ensure:

109 Support for differentiated services and different traffic types with traffic class separation associated with different traffic.

110 Enabling the provisioning and the guarantee of Service Level Specifications (SLS), with support for hard and relative end-to-end bandwidth guaranteed.

111 Support of services, which are sensitive to jitter and delay.

112 Guarantee of fair access, within a particular class, to shared resources.

113 Guaranteed resources for in-band control and management plane traffic regardless of the amount of data plane traffic.

114 Carriers are provided with the capability to efficiently support service demands over the MPLS-TP network. This MUST include support for a flexible bandwidth allocation scheme.

<<Comment #33>>
Please add the following requirement:
"

It MUST be possible to support operator's policies regarding agreed traffic profiles.
"
<<End Comment>>

2.10. Security requirements

For a description of the security threats relevant in the context of MPLS and GMPLS and the defensive techniques to combat those threats see the Security Framework for MPLS & GMPLS Networks [I-D.ietf-mpls-mpls-and-gmpls-security-framework].

3. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

4. Security Considerations

See Section 2.10.

5. Acknowledgements

The authors would like to thank all members of the teams (the Joint Working Team, the MPLS Interoperability Design Team in the IETF, and the T-MPLS Ad Hoc Group in the ITU-T) involved in the definition and specification of MPLS Transport Profile.

The authors would also like to thank Loa Andersson, Dieter Beller, Lou Berger, Italo Busi, John Drake, Adrian Farrel, Annamaria Fulignoli, Pietro Grandi, Eric Gray, Neil Harrison, Huub van Helvoort, Enrique Hernandez-Valencia, Wataru Imajuku, Kam Lam, Andy Malis, Alan McGuire, Julien Meuric, Greg Mirsky, Tom Nadeau, Hiroshi Ohta, Tom Petch, Andy Reid, Vincenzo Sestito, George Swallow, Lubo Tancevski, Tomonori Takeda, Yuji Tochio, Alexander Vainshtein, Eve Varma and Maarten Vissers for their comments and enhancements to the text.

An ad hoc discussion group consisting of Stewart Bryant, Italo Busi, Andrea Digiglio, Li Fang, Adrian Farrel, Jia He, Huub van Helvoort, Feng Huang, Harald Kullman, Han Li, Hao Long and Nurit Sprecher provided valuable input to the requirements for deployment and survivability in ring topologies.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-

Edge (PWE3) Architecture", RFC 3985, March 2005.

[RFC4427] Mannie, E. and D. Papadimitriou, "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, March 2006.

[I-D.ietf-mpls-mpls-and-gmpls-security-framework] Fang, L. and M. Behringer, "Security Framework for MPLS and GMPLS Networks", draft-ietf-mpls-mpls-and-gmpls-security-framework-05 (work in progress), November 2008.

[I-D.ietf-pwe3-ms-pw-arch] Bocci, M. and S. Bryant, "Requirements for OAM in MPLS Transport Networks", draft-ietf-pwe3-ms-pw-arch-06 (work in progress), September 2008.

[ITU.G805.2000] International Telecommunications Union, "Generic functional architecture of transport networks", ITU-T Recommendation G.805, March 2000.

[ITU.G8080.2006] International Telecommunications Union, "Architecture for the automatically switched optical network (ASON)", ITU-T Recommendation G.8080, June 2006.

[ITU.G8080.2008] International Telecommunications Union, "Architecture for the automatically switched optical network (ASON) Amendment 1", ITU-T Recommendation G.8080 Amendment 1, March 2008.

6.2. Informative References

[RFC4139] Papadimitriou, D., Drake, J., Ash, J., Farrel, A., and L. Ong, "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)", RFC 4139, July 2005.

[RFC4258] Brungard, D., "Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Routing for the Automatically Switched Optical Network (ASON)", RFC 4258, November 2005.

[RFC4397] Bryskin, I. and A. Farrel, "A Lexicography for the Interpretation of Generalized Multiprotocol Label Switching (GMPLS) Terminology within the Context of the ITU-T's Automatically Switched Optical Network (ASON)

Architecture", RFC 4397, February 2006.

[I-D.ietf-mpls-tp-nm-req] Lam, H., Mansfield, S., and E. Gray, "MPLS TP Network Management Requirements", draft-ietf-mpls-tp-nm-req-00 (work in progress), January 2009.

[I-D.ietf-mpls-tp-oam-requirements] Vigoureux, M., Ward, D., and M. Betts, "Requirements for OAM in MPLS Transport Networks", draft-ietf-mpls-tp-oam-requirements-01 (work in progress), November 2008.

[ITU.Y1401.2008] International Telecommunications Union, "Principles of interworking", ITU-T Recommendation Y.1401, February 2008.

[ITU.Y2611.2006] International Telecommunications Union, "High-level architecture of future packet-based networks", ITU-T Recommendation Y.2611, December 2006.

Authors' Addresses

Ben Niven-Jenkins (editor)
BT
208 Callisto House, Adastral Park
Ipswich, Suffolk IP5 3RE
UK

Email: benjamin.niven-jenkins@bt.com

Deborah Brungard (editor)
AT&T
Rm. D1-3C22 - 200 S. Laurel Ave.
Middletown, NJ 07748
USA

Email: dbrungard@att.com

Malcolm Betts (editor)
Nortel Networks
3500 Carling Avenue
Ottawa, Ontario K2H 8E9
Canada

Email: betts01@nortel.com

Nurit Sprecher
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon,   45241
Israel

Email: nurit.sprecher@nsn.com

Satoshi Ueno
NTT

Email: satoshi.ueno@ntt.com